arXiv:1410.2726v1 [math.OC] 10 Oct 2014

Constrained Stochastic Optimal Control with a Baseline Performance Guarantee

Yinlam Chow
Institute of Computational & Mathematical Engineering, Stanford University
Mohammad Ghavamzadeh*
INRIA Lille - Team SequeL & Adobe Research

Abstract

In this paper, we show how a simulated Markov decision process (MDP) built by the so-called *baseline* policies, can be used to compute a different policy, namely the *simulated optimal* policy, for which the performance of this policy is guaranteed to be better than the baseline policy in the real environment. This technique has immense applications in fields such as news recommendation systems, health care diagnosis and digital online marketing. Our proposed algorithm iteratively solves for a "good" policy in the simulated MDP in an offline setting. Furthermore, we provide a performance bound on sub-optimality for the control policy generated by the proposed algorithm.

1 Introduction

In this paper, we show how a simulated Markov decision process (MDP) built by the so-called *baseline* policies, can be used to compute a different policy, namely the *simulated optimal* policy, for which the performance of this policy is guaranteed to be better than the baseline policy in the real environment. This policy evaluation technique has immense applications in various fields, where offline calculation of a new policy is computationally inexpensive, but execution of a new policy can be hazardous or costly if this new policy performs worse than the baseline policy. There are numerous applications that require such safety concerns. For example, in the personalized news article recommendation systems [8], one requires a "good" policy that selects articles to serve users based on contextual information in order to maximize total user clicks. In this online setup, real-time policy evaluation is dangerous in the sense that a slight fluctuation of performance may yield a huge loss in click rates. Similar examples can be found in patient diagnosis systems [6, 4] as well, where testing new clinical strategies on human trials is often too risky. Further applications on medical trials can be found in automatic Propofol administration [13] and neuroprosthetic control [9]. Last but not least, under the framework of digital online marketing [7], one also require a "good" ad-recommendation policy that sequentially displays attractive ads in order to increase customers' click-rates while the performance of this policy is guaranteed to be better than the company's baseline marketing strategy.

In this paper, we make several assumptions. First, we assume the evolution of the real environment follows a Markov decision process. This assumption is standard in most literatures of sequential decision making problems [3, 11]. Second, we also assume that a deterministic error upper bound function between the true MDP and the simulated MDP is known. While this assumption can be justified probabilistically using the Chernoff-hoeffding inequality when the state and action spaces are finite, one may find it challenging to obtain such an error bound in advance for systems with continuous state and action spaces, especially when there is no prior knowledge in the real system model. Nevertheless, from the best of our knowledge, the method proposed in this paper is still

*Mohammad Ghavamzadeh is at Adobe Research, on leave of absence from INRIA Lille - Team SequeL.

novel and applicable in many areas. One can easily extend the analysis of our methods to include the case of probabilistic error bounds as well.

Since our problem formulation guarantees that the resultant control policy performs better than its baseline counterpart, one may connect our method to the vast literature of stochastic control problems with system uncertainties [10, 14, 12]. The latter approach solves for a robust policy which guarantees worst-case performance. While our proposed method can also be posed as a stochastic control problem with system uncertainties characterized by the error upper bound, the resultant policy may be over-conservative. Furthermore, we are only interested in its performance under the real environment, rather than the worst case performance. On top of that, in the worst-case approach, various assumptions are required for the set of system uncertainties in order to construct numerically tractable algorithms. For example, see the rectangularity assumption of uncertain transition probability set in [5]. Therefore, our proposed algorithm in general provides a simple yet better control policy, compared to its robust counterpart.

The rest of this paper is structured as follows. In Section 2, preliminary notations and definitions used in this paper are discussed. The constrained MDP problem formulation is given in Section 3 where further analysis on the Lagrangian formulation based on state augmentations is given in Section 4. Based on the strong duality result in Section 5, one can equivalently consider the dual Lagrangian formulation. On top of this, by using the feasibility condition in Section 6, an iterative algorithm is provided in Section 7, followed by a performance analysis of sub-optimality in Section 8. Finally the results of this paper is concluded in Section 9.

2 Preliminaries

We consider problems in which the agent's interaction with the environment is modeled as a MDP. A MDP is a tuple $M = (\mathcal{X}, \mathcal{A}, R, P, P_0)$, where $\mathcal{X} = \{1, \ldots, n\}$ and $\mathcal{A} = \{1, \ldots, m\}$ are the state and action spaces; $R(x, a) \in [-R_{\max}, R_{\max}]$ is the bounded cost random variable whose expectation is denoted by $r(x, a) = \mathbb{E}\big[R(x, a)\big]$; $P(\cdot|x, a)$ is the transition probability distribution; and $P_0(\cdot)$ is the initial state distribution. For simplicity, we assume that the system has a single initial state x^0, i.e., $P_0(x) = \mathbf{1}\{x = x^0\}$. All the results of the paper can be easily extended to the case that the system has more than one initial state. We also need to specify the rule according to which the agent selects actions at each state. A *stationary policy* $\mu(\cdot|x)$ is a probability distribution over actions, conditioned on the current state. We denote by $d_\gamma^\mu(x|x^0) = (1-\gamma)\sum_{k=0}^\infty \gamma^k \mathbb{P}(x_k = x|x_0 = x^0; \mu)$ and $\pi_\gamma^\mu(x, a|x^0) = d_\gamma^\mu(x|x^0)\mu(a|x)$ the γ-discounted visiting distribution of state x and state-action pair (x, a) under policy μ, respectively.

However, even if the state/action spaces and reward functions are given, in many cases the real world model M is unknown, in particular, the underlying transition probability P cannot be easily obtained by the user. Rather, a simulated transition probability $\widehat{P}$ is often obtained by Monte Carlo sampling techniques with a baseline policy distribution μ_B. In many engineering applications, the following simulated MDP model $\widehat{M} = (\mathcal{X}, \mathcal{A}, R, \widehat{P}, P_0)$ can be built based on previous histories of data and Monte Carlo sampling techniques. There are many approaches to search for a "good" control policy using the information of MDP model $\widehat{M}$. As the most direct approach, one can search for an optimal policy over $\widehat{M}$ with no interactions with the real world using dynamic programming or approximate dynamic programming methods. Unfortunately, while it could be computationally expensive to find an optimal policy from $\widehat{M}$ when the state/action spaces are huge, this optimal policy in general do not have a performance guarantee to the original model M.

In order to characterize the error between the true and simulated models, define the "mis-measure" of transition probability at $(x, a) \in \mathcal{X} \times \mathcal{A}$ as the 1-norm of the difference between P and $\widehat{P}$, i.e.,

$$\|P(\cdot|x, a) - \widehat{P}(\cdot|x, a)\|_1 = \sum_{y \in \mathcal{X}} \Big| P(y|x, a) - \widehat{P}(y|x, a) \Big| .$$

Suppose an upper bound $e(x, a)$ of the mis-measure is given, i.e.,

$$e(x, a) \geq \sum_{y \in \mathcal{X}} \Big| P(y|x, a) - \widehat{P}(y|x, a) \Big| .$$

Using the notion of transition probability "mis-measure" , we will later illustrate a policy search method in this paper that finds an optimal control policy in the simulated model with guaranteed performance improvements over the baseline policy μ_B.

Before getting into the main results, we first introduce several policy spaces. Define the space H_t of admissible histories up to time t by $H_t = H_{t-1} \times \mathcal{X} \times \mathcal{A}$, for $t \geq 1$, and $H_0 = \mathcal{X}$. A generic element $h_t \in H_t$ is of the form $h_t = (x_0, a_0, \ldots, x_{t-1}, a_{t-1}, x_t)$. Let Π_t be the set of all stochastic history dependent policies with the property that at each time t the control is a function of h_t. In other words, $\Pi_{H,t} := \Big\{\{\mu_0 : H_0 \to \mathbb{P}(\mathcal{A}), \mu_1 : H_1 \to \mathbb{P}(\mathcal{A}), \ldots, \mu_{t-1} : H_{t-1} \to \mathbb{P}(\mathcal{A})\} | \mu_k(h_k) \in \mathbb{P}(\mathcal{A}) \text{ for all } h_k \in H_k,\, t-1 \geq k \geq 0\Big\}$. While history dependent policy space is the most general space for control policies that satisfy the causality assumption, it is often times computationally intractable to search for an optimal policy over this space. On the other hand, we define the set of Markov policies as a sequence of state-action mappings as follows: $\Pi_{M,t} := \Big\{\{\mu_0 : \mathcal{X} \to \mathbb{P}(\mathcal{A}), \mu_1 : \mathcal{X} \to \mathbb{P}(\mathcal{A}), \ldots, \mu_{t-1} : \mathcal{X} \to \mathbb{P}(\mathcal{A})\} | \mu_k(x_k) \in \mathbb{P}(\mathcal{A}) \text{ for all } x_k \in \mathcal{X},\, t-1 \geq k \geq 0\Big\}$. Furthermore, when the state-action mapping is stationary across time, we define the space of stationary Markov policies as follows $\Pi_{S,t} := \Big\{\{\mu : \mathcal{X} \to \mathbb{P}(\mathcal{A}), \mu : \mathcal{X} \to \mathbb{P}(\mathcal{A}), \ldots, \mu : \mathcal{X} \to \mathbb{P}(\mathcal{A})\} | \mu(x_k) \in \mathbb{P}(\mathcal{A}) \text{ for all } x_k \in \mathcal{X},\, t-1 \geq k \geq 0\Big\}$. We also let $\Pi_H = \lim_{t\to\infty} \Pi_{H,t}$, $\Pi_M = \lim_{t\to\infty} \Pi_{M,t}$ and $\Pi_S = \lim_{t\to\infty} \Pi_{S,t}$ be the set of history dependent policies, Markov policies and stationary Markov policies respectively. We will later show that the proposed method performs policy search over the space of stationary Markov policies without loss of optimality.

3 Problem Formulation

We propose the following stochastic optimal control problem whose optimal policy $\mathcal{U}^* = \{\mu_0^*, \mu_1^*, \ldots\}$, 1) maximizes the expected return in model based MDP $\widehat{M}$ and 2) guarantees a better performance then $\mathcal{U}^B = \{\mu_0^B, \mu_1^B, \ldots\}$ in the real world MDP M,

$$\begin{aligned} \max_{\mathcal{U}\in\Pi_H} \quad & \mathbb{E}_{\widehat{P}}\left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu_k, P_0\right) \\ \text{subject to} \quad & \mathbb{E}_{P}\left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu_k, P_0\right) \geq \mathbb{E}_{P}\left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu_k^B, P_0\right). \end{aligned}$$

Now define the Lagrangian function as

$$\begin{aligned} L(\mathcal{U}, \lambda) \stackrel{\triangle}{=} & \mathbb{E}_{\widehat{P}}\left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu_k, P_0\right) \\ & + \lambda\left(\mathbb{E}_{P}\left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu_k, P_0\right) - \mathbb{E}_{P}\left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu_k^B, P_0\right)\right). \end{aligned}$$

Since the baseline policy $\mathcal{U}^B$ is given, we define

$$M_B = \mathbb{E}_{P}\left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu_k^B, P_0\right)$$

as the safety threshold for the above problem and assume this quantity is known to the user. To solve the constrained risk sensitive optimization problem, we employ the Lagrangian formulation to convert the above problem into the following unconstrained problem:

$$\max_{\mathcal{U}\in\Pi_H} \min_{\lambda\geq 0} L(\mathcal{U}, \lambda). \tag{1}$$

4 The Augmented MDP

Since the Lagrangian function involves multiple probability distributions, the above problem cannot be easily solved using existing stochastic control methods from the MDP literature. Therefore, one constructs the following augmented MDP in order to rewrite $L(\mu, \lambda)$ under one MDP. Consider an augmented MDP $M_\lambda^{\text{aug}} = (\bar{\mathcal{X}}, \bar{\mathcal{A}}, P^{\text{aug}}, r_\lambda, \gamma, P_0^{\text{aug}})$ such that $\bar{\mathcal{X}} = \mathcal{X} \times I$, where $I = \{0, 1\}$, $\bar{\mathcal{A}} = \mathcal{A}$, γ is the discounted factor given above and,

$$r_\lambda(x, i, a) = \begin{cases} 2r(x, a)\mathbf{1}\{x = x_T\} & \text{if } i = 0 \\ 2\lambda r(x, a)\mathbf{1}\{x = x_T\} & \text{if } i = 1 \end{cases}, \ \forall (x, i) \in \bar{\mathcal{X}}, \ a \in \bar{\mathcal{A}},$$

$$P_0^{\text{aug}}(x, i) = \begin{cases} \frac{1}{2}P_0(x) & \text{if } i = 0 \\ \frac{1}{2}P_0(x) & \text{if } i = 1 \end{cases}, \ \forall (x, i) \in \bar{\mathcal{X}}.$$

We also define the uncertain transition probability as follows:

$$P^{\text{aug}}(x', i'|x, i, a) = \begin{cases} \widehat{P}(x'|x, a)\mathbf{1}\{i' = 0\} & \text{if } i = 0 \\ P(x'|x, a)\mathbf{1}\{i' = 1\} & \text{if } i = 1 \end{cases}, \ \forall (x', i'), (x, i) \in \bar{\mathcal{X}}.$$

For any $(x, i) \in \bar{\mathcal{X}}$, define the robust $\lambda-$parametrized optimal Bellman's operator $T_\lambda : \mathbb{R}^{|\bar{\mathcal{X}}|} \to \mathbb{R}^{|\bar{\mathcal{X}}|}$ as follows:

$$T_\lambda[V](x, i) = \max_{a \in \bar{\mathcal{A}}} \left\{ r_\lambda(x, i, a) + \gamma \sum_{(x', i') \in \bar{\mathcal{X}}} P^{\text{aug}}(x', i'|x, i, a) V(x', i') \right\}.$$

In the following analysis, we will assume $\lambda \geq 0$ is bounded[1]. From the literature in dynamic programming, one can easily see that the above Bellman operator satisfies the following properties, and the fixed point solution converges to the optimal solution of the Lagrangian function. For the proofs of these properties, please see Chapter 1 of [3] for more details.

Proposition 1. *The Bellman operator $T_\lambda[V]$ has the following properties:*

- *(Monotonicity) If $V_1(x, i) \geq V_2(x, i)$, for any $x \in \mathcal{X}$, $i \in I$ then $T_\lambda[V_1](x, i) \geq T_\lambda[V_2](x, i)$.*
- *(Constant shift) For $K \in \mathbb{R}$, $T_\lambda[V + K](x, i) = T_\lambda[V](x, i) + \gamma K$.*
- *(Contraction) For any $V_1, V_2 : \mathcal{X} \times I \to \mathbb{R}$, $\|T_\lambda[V_1] - T_\lambda[V_2]\|_\infty \leq \gamma \|V_1 - V_2\|_\infty$, where $\|f\|_\infty = \max_{x \in \mathcal{X}, i \in I} |f(x, i)|$.*

Theorem 2. *For any $\lambda \geq 0$, there exists a unique solution to the fixed point equation: $T_\lambda[V](x, i) = V(x, i)$, $\forall (x, i) \in \bar{\mathcal{X}}$. Let $V_\lambda \in \mathbb{R}^{|\bar{\mathcal{X}}|}$ be such unique fixed point solution. For any initial value function $V_{\lambda,0}(x, i)$, $V_\lambda(x, i) = \lim_{N \to \infty} T_\lambda^N[V][V_{\lambda,0}](x, i)$, $\forall (x, i) \in \bar{\mathcal{X}}$. Furthermore,*

$$\sum_{x,i} P_0^{aug}(x, i) V_\lambda(x, i) = \max_{\mathcal{U} \in \Pi_H} L(\mathcal{U}, \lambda) + \lambda M_B. \tag{2}$$

By the above theory of Bellman's recursion, one can show that the space of augmented state feedback Markov policies is dominating.

Theorem 3 (Dominating Policies)**.** *Suppose V_λ is the unique solution to $T_\lambda[V](x, i) = V(x, i)$, $\forall (x, i) \in \bar{\mathcal{X}}$. Then, suppose the policy μ is found by*

$$\mu(x, i) \in \arg\max_{a \in \bar{\mathcal{A}}} \left\{ r_\lambda(x, i, a) + \gamma \sum_{(x', i') \in \bar{\mathcal{X}}} P^{aug}(x', i'|x, i, a) V(x', i') \right\}.$$

Then for $\mathcal{U}^* = \{\mu, \mu, \ldots\}$ being a sequence of Markov stationary optimal policies, it follows that $\mathcal{U}^* \in \arg\max_{\mathcal{U} \in \Pi_H} L(\mathcal{U}, \lambda) + \lambda M_B$. In other words, the class of $\lambda-$parametrized Markov stationary deterministic policies is dominating.

[1] The boundedness assumption is justified in the solution algorithm section.

Proof. Suppose $T_\lambda^\mu[V_\lambda](x,i) = T_\lambda[V_\lambda](x,i)$ for any $(x,i) \in \bar{\mathcal{X}}$, where the Bellman operator $T_\lambda^\mu[V]$ is defined as follows:

$$T_\lambda^\mu[V](x,i) = r_\lambda(x,i,\mu(x,i)) + \gamma \sum_{(x',i')\in\bar{\mathcal{X}}} P^{\text{aug}}(x',i'|x,i,\mu(x,i))V(x',i').$$

Then, from the above equality and the fixed point equation, $T_\lambda^\mu[V_\lambda](x,i) = T_\lambda[V_\lambda](x,i) = V_\lambda(x,i)$. By the unique-ness of the fixed point equation: $T_\lambda^\mu[V](x,i) = V(x,i)$, one also obtains $V_\lambda(x,i) = V_\lambda^\mu(x,i)$, where $V_\lambda^\mu(x,i) = \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^\infty \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, x_0 = x, i_0 = i\right)$. By summing over $(x,i) \in \bar{\mathcal{X}}$, weighted by the initial distribution $\{P_0^{\text{aug}}(x,i)\}_{(x,i)\in\bar{\mathcal{X}}}$, one obtains

$$\begin{aligned}\max_{\mathcal{U}\in\Pi_H} L(\mathcal{U},\lambda) + \lambda M_B =& \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^\infty \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, P_0^{\text{aug}}\right)\\ =& \sum_{x,i} P_0^{\text{aug}}(x,i)V_\lambda^\mu(x,i) = \sum_{x,i} P_0^{\text{aug}}(x,i)V_\lambda(x,i),\end{aligned}$$

which further implies μ^* is a sequence of Markov stationary optimal policies (and μ is optimal). □

Since the class of stationary Markov policies is dominating, without loss of generality, one can write

$$\max_{\mathcal{U}\in\Pi_H}\min_{\lambda\geq 0} L(\mathcal{U},\lambda) = \max_{\mu\in\Pi_S}\min_{\lambda\geq 0} L(\mu,\lambda),$$

where

$$L(\mu,\lambda) = \mathbb{E}_{\widehat{P}}\left(\sum_{k=0}^\infty \gamma^k r(x_k,a_k) \mid a_k \sim \mu, P_0\right) + \lambda\mathbb{E}_P\left(\sum_{k=0}^\infty \gamma^k r(x_k,a_k) \mid a_k \sim \mu, P_0\right) - \lambda M_B.$$

and μ belongs to the set of stationary Markov policies of the augmented MDP M_λ^{aug}.

5 Strong Duality

In previous sections, one transforms the constrained optimization problem into its primal Lagrangian formulation. However in often times, the dual Lagrangian formulation often yields a computationally traceable and efficient algorithm. In this section, we want to show that strong duality exists, thus one can equivalently switch from the primal formulation to its dual counterpart. Consider the primal optimization problem:

$$\mathcal{P}(P_0^{\text{aug}}, M_B) = \max_\mu \min_{\lambda\geq 0} \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^\infty \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, P_0^{\text{aug}}\right) - \lambda M_B = \max_\mu \min_{\lambda\geq 0} L(\mu,\lambda),$$

and the dual problem:

$$\mathcal{D}(P_0^{\text{aug}}, M_B) = \min_{\lambda\geq 0} \max_\mu \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^\infty \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, P_0^{\text{aug}}\right) - \lambda M_B = \min_{\lambda\geq 0} \max_\mu L(\mu,\lambda).$$

Now consider the risk neutral optimization problem. Recall from that previous section that

$$\max_\mu \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^\infty \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, P_0^{\text{aug}}\right) = \sum_{(x,i)\in\bar{\mathcal{X}}} P_0^{\text{aug}}(x,i)V_\lambda(x,i),$$

where $V_\lambda(x,i)$ is the unique solution of the following Bellman equation $T_\lambda[V](x,i) = V(x,i)$. Furthermore, by the linear programming formulation of Bellman equation, one obtains

$$\begin{aligned}\text{RLP}(\lambda) = \min_V \quad & \sum_{(x,i)\in\bar{\mathcal{X}}} P_0^{\text{aug}}(x,i)V(x,i) && (3)\\ \text{s.t.} \quad & r_\lambda(x,i,a) + \gamma \sum_{(x',i')\in\bar{\mathcal{X}}} P^{\text{aug}}(x',i'|x,i,a)V(x',i') \leq V(x,i),\ \forall (x,i)\in\bar{\mathcal{X}},\ a\in\bar{\mathcal{A}}\\ = \max_\mu \quad & \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^\infty \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, P_0^{\text{aug}}\right). && (4)\end{aligned}$$

Consider the Lagrangian of the above problem with Lagrangian variable $\{\beta(x,i,a)\}_{(x,i)\in\bar{\mathcal{X}},a\in\bar{\mathcal{A}}} \in \mathbb{R}_{\geq 0}^{|\bar{\mathcal{X}}|\times|\bar{\mathcal{A}}|}$:

$$\mathcal{L}_\lambda(V,\beta) = \sum_{(x,i)\in\bar{\mathcal{X}}} P_0^{\text{aug}}(x,i)V(x,i) + \sum_{(x,i)\in\bar{\mathcal{X}},a\in\bar{\mathcal{A}}} \beta(x,i,a)\cdot \left\{ r_\lambda(x,i,a) + \gamma \sum_{x'\in\mathcal{X}} P^{\text{aug}}(x',i'|x,i,a)V(x',s',i') - V(x,i) \right\},$$

where

$$\min_V \max_{\beta\in\mathbb{R}_{\geq 0}^{|\bar{\mathcal{X}}|\times|\bar{\mathcal{A}}|}} \mathcal{L}_\lambda(V,\beta) = \max_\mu \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty} \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, P_0^{\text{aug}} \right) = \text{RLP}(\lambda)$$

follows from the primal argument of Lagrangian duality theory. Note that $\mathcal{L}_\lambda(V,\beta)$ is a concave function in β and convex function in V. Assume Slater's condition holds, this implies the following result:

$$\min_V \max_{\beta\in\mathbb{R}_{\geq 0}^{|\bar{\mathcal{X}}|\times|\bar{\mathcal{A}}|}} \mathcal{L}_\lambda(V,\beta) = \max_{\beta\in\mathbb{R}_{\geq 0}^{|\bar{\mathcal{X}}|\times|\bar{\mathcal{A}}|}} \min_V \mathcal{L}_\lambda(V,\beta).$$

Before proving the main result, we first state the Aubin's Minmax theorem.

Lemma 4 (Minmax Theorem [2])**.** *Let G_1 and G_2 be convex subsets of linear topological spaces, and let G_1 be compact. Consider a function $\psi : G_1 \times G_2 \to \mathbb{R}$ such that*

- *for each $g_2 \in G_2$, $g_1 \to \psi(g_1,g_2)$ is convex and lower semi-continuous, and*
- *for each $g_1 \in G_1$, $g_2 \to \psi(g_1,g_2)$ is concave.*

Then there exists some $g_1^* \in G_1$ such that

$$\inf_{G_1}\sup_{G_2} \psi(g_1,g_2) = \sup_{G_2} \psi(g_1^*,g_2) = \sup_{G_2}\inf_{G_1} \psi(g_1,g_2).$$

Then, the next theorem shows that strong duality holds.

Theorem 5. *For any given initial distribution P_0^{aug} and safety threshold M_B, the following statement holds:*

$$\mathcal{D}(P_0^{aug}, M_B) = \mathcal{P}(P_0^{aug}, M_B).$$

Proof. First consider the following expression:

$$\begin{aligned} &\min_{\lambda\geq 0}\max_\mu \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty} \gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k \sim \mu, P_0^{\text{aug}} \right) \\ =& \min_{V,\lambda\geq 0} \max_{\beta\in\mathbb{R}_{\geq 0}^{|\bar{\mathcal{X}}|\times|\bar{\mathcal{A}}|}} \mathcal{L}_\lambda(V,\beta) = \min_{\lambda\geq 0} \max_{\beta\in\mathbb{R}_{\geq 0}^{|\bar{\mathcal{X}}|\times|\bar{\mathcal{A}}|}} \min_V \mathcal{L}_\lambda(V,\beta). \end{aligned}$$

The first equality follows from the previous primal arguments of duality theory and the second equality follows from strong duality. Also, consider the following function

$$\Psi(\lambda,\beta) := \min_V \mathcal{L}_\lambda(V,\beta) = \sum_{(x,i)\in\bar{\mathcal{X}},a\in\bar{\mathcal{A}}} \beta(x,i,a) r_\lambda(x,i,a) + \Phi(\beta),$$

where

$$\begin{aligned} \Phi(\beta) = \min_V &\sum_{x\in\mathcal{X}} P_0^{\text{aug}}(x,0)V(x,0) + \sum_{x\in\mathcal{X},a\in\bar{\mathcal{A}}} \beta(x,0,a)\left\{ \gamma \sum_{x'\in\mathcal{X}} \widehat{P}(x'|x,a)V(x',0) - V(x,0) \right\} \\ &+ \sum_{x\in\mathcal{X}} P_0^{\text{aug}}(x,1)V(x,1) + \sum_{x\in\mathcal{X},a\in\bar{\mathcal{A}}} \beta(x,1,a)\left\{ \gamma \sum_{x'\in\mathcal{X}} P(x'|x,a)V(x',1) - V(x,1) \right\}. \end{aligned}$$

This implies that whenever $\Phi(\beta)$ is bounded, the following condition holds:

$$P_0^{\text{aug}}(x,1) = \sum_{a\in\bar{\mathcal{A}}} \beta(x,1,a) \sum_{x'\in\mathcal{X}} \delta\{x=x'\} - \gamma P(x'|x,a),\ \forall x\in\mathcal{X}, \tag{5}$$

$$P_0^{\text{aug}}(x,0) = \sum_{a\in\bar{\mathcal{A}}} \beta(x,0,a) \sum_{x'\in\mathcal{X}} \delta\{x=x'\} - \gamma \widehat{P}(x'|x,a),\ \forall x\in\mathcal{X}. \tag{6}$$

Then, one concludes that

$$\Phi(\beta) = \begin{cases} 0 & \text{if } P_0^{\text{aug}}(x,i) = \sum_{a\in\bar{\mathcal{A}}} \beta(x,i,a) \sum_{x',i'} \delta\{x=x',i=i'\} - \gamma P^{\text{aug}}(x',i'|x,i,a),\ \forall x,i \\ -\infty & \text{otherwise} \end{cases}$$

and

$$\Psi(\lambda,\beta) = \begin{cases} \sum_{x,i,a} \beta(x,i,a) r_\lambda(x,i,a)/(1-\gamma) & \text{if } \beta\in\mathcal{B} \\ -\infty & \text{otherwise} \end{cases}$$

where

$$\mathcal{B} = \left\{ \beta\in\mathbb{R}_{\geq 0}^{|\bar{\mathcal{X}}|\times|\bar{\mathcal{A}}|} : (1-\gamma)P_0^{\text{aug}}(x,i) = \sum_{a\in\bar{\mathcal{A}}} \beta(x,i,a) \sum_{x',i'} \delta\{x=x',i=i'\} - \gamma P^{\text{aug}}(x',i'|x,i,a),\ \forall x,i \right\}.$$

Note that for any $\beta\in\mathcal{B}$, $\sum_{x,i,a}\beta(x,i,a)=1$. Since $\Psi(\lambda,\beta)-\lambda M_B$ is a linear function in λ, it is also a convex and lower semi-continuous function in λ. Furthermore $\{\lambda\geq 0\}$ is a convex compact set of λ. On the other hand, $\Psi(\lambda,\beta)$ is a concave function in β, due to the facts that 1) $\sum_{(x,i)\in\bar{\mathcal{X}},a\in\bar{\mathcal{A}}}\beta(x,i,a)r_\lambda(x,i,a)/(1-\gamma)$ is a linear function of β and 2) $\mathcal{B}$ is a convex set. Therefore, by Aubin's Minimax Theorem, one concludes that

$$\begin{aligned}\min_{\lambda\geq 0}\max_{\beta\in\mathcal{B}}\min_{V} \mathcal{L}_\lambda(V,\beta) - \lambda M_B &= \min_{\lambda\geq 0}\max_{\beta\in\mathcal{B}} \Psi(\lambda,\beta) - \lambda M_B \\ &= \max_{\beta\in\mathcal{B}}\min_{\lambda\geq 0} \Psi(\lambda,\beta) - \lambda M_B = \max_{\beta\in\mathcal{B}}\min_{\lambda\geq 0}\min_{V} \mathcal{L}_\lambda(V,\beta) - \lambda M_B.\end{aligned}$$

The final step is to prove

$$\max_{\beta\in\mathcal{B}}\min_{\lambda\geq 0} \sum_{(x,i)\in\bar{\mathcal{X}},a\in\bar{\mathcal{A}}} \beta(x,i,a)\frac{r_\lambda(x,i,a)}{1-\gamma} - \lambda M_B = \max_{\mu}\min_{\lambda\geq 0} \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k\sim\mu, P_0^{\text{aug}}\right) - \lambda M_B.$$

The proof of this statement follows directly from Theorem 3.1 in [1]. □

We have just shown that strong duality holds in this optimization problem. Thus, in the following sections, we will mainly focus on deriving algorithms to solve the dual problem described in $\mathcal{D}(P_0^{\text{aug}}, M_B)$.

6 A Condition for Guaranteeing Feasibility

Recall the risk neutral maximization problem is equivalent to its primal Lagrangian formulation

$$\max_{\mu}\min_{\lambda\geq 0} L(\mu,\lambda).$$

However when the MDP M is unknown, one cannot directly solve this problem. Therefore, we want to derive a sufficient condition to guarantee feasibility based on solving a similar problem using MDP $\widehat{M}$. Note that the risk neutral maximization problem is feasible if $\max_\mu \min_{\lambda\geq 0} L(\mu,\lambda)$ is lower bounded. Therefore if there exists a function $\widehat{L}(\mu,\lambda)$ such that $\widehat{L}(\mu,\lambda)\leq L(\mu,\lambda)$, then

$$\min_{\lambda\geq 0} L(\mu,\lambda) := L(\mu,\lambda^*) \geq \widehat{L}(\mu,\lambda^*) \geq \min_{\lambda\geq 0}\widehat{L}(\mu,\lambda).$$

By taking maximization over μ on both sides, one further obtains

$$\max_{\mu}\min_{\lambda\geq 0} L(\mu,\lambda) \geq \max_{\mu}\min_{\lambda\geq 0}\widehat{L}(\mu,\lambda).$$

Therefore, if $\max_\mu\min_{\lambda\geq 0}\widehat{L}(\mu,\lambda)$ is lower bounded, one concludes that the original problem is always feasible. Now, consider the following construction of $\widehat{L}(\mu,\lambda)$.

Lemma 6. *Define*

$$\widehat{L}(\mu,\lambda) := \mathbb{E}_{\widehat{P}^{aug}}\left(\sum_{k=0}^{\infty}\gamma^k \widehat{r}_\lambda(x_k,i_k,a_k) \mid a_k\sim\mu, P_0^{aug}\right) - \lambda M_B,$$

where

$$\widehat{r}_\lambda(x,i,a) = \begin{cases} 2r(x,a) & \text{if } i=0 \\ 2\lambda r(x,a) - \frac{2\gamma\lambda R_{\max}}{1-\gamma}e(x,a) & \text{if } i=1 \end{cases}, \ \forall (x,i)\in\bar{\mathcal{X}},\ a\in\bar{\mathcal{A}}.$$

Notice that $\widehat{P}^{aug}$ is the transition probability in the augmented state and action spaces such that its definition is similar to the definition of P^{aug}, except by replacing P by $\widehat{P}$.

Then for any stationary Markovian policy μ and Lagrangian multiplier $\lambda \geq 0$, the following inequality holds:

$$\begin{aligned}
&L(\mu,\lambda) - \widehat{L}(\mu,\lambda) \\
=&\mathbb{E}_{P^{aug}}\left(\sum_{k=0}^{\infty}\gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k\sim\mu, P_0^{aug}\right) - \mathbb{E}_{\widehat{P}^{aug}}\left(\sum_{k=0}^{\infty}\gamma^k \widehat{r}_\lambda(x_k,i_k,a_k) \mid a_k\sim\mu, P_0^{aug}\right) \geq 0.
\end{aligned}$$

Proof. For any transition probability P^{aug} and any Markov policy μ of the augmented MDP, define

$$d^{\mu}_{P^{\text{aug}},\gamma}(x,i) = \sum_{k=0}^{\infty}\gamma^k(1-\gamma)\sum_{a\in\mathcal{A}}\mathbb{P}_{P^{\text{aug}}}(x_k=x, i_k=i | x_0=x^0, i_0=i^0,\mu).$$

where the occupation measure satisfies the following expression:

$$\begin{aligned}
&\gamma \sum_{(x',i')\in\bar{\mathcal{X}}} d^{\mu}_{P^{\text{aug}},\gamma}(x',i') \sum_{a'\in\bar{\mathcal{A}}} P^{\text{aug}}(x,i|x',i',a')\mu(a'|x',i') \\
=&(1-\gamma)\sum_{(x',i')\in\bar{\mathcal{X}}}\sum_{a'\in\bar{\mathcal{A}}}\sum_{k=0}^{\infty}\gamma^{k+1}\mathbb{P}_{P^{\text{aug}}}(x_k=x', i_k=i' | x_0=x^0, i_0=i^0,\mu)P^{\text{aug}}(x,i|x',i',a')\mu(a'|x',i') \\
=&(1-\gamma)\sum_{k=0}^{\infty}\gamma^{k+1}\mathbb{P}_{P^{\text{aug}}}(x_{k+1}=x, i_{k+1}=i | x_0=x^0, i_0=i^0,\mu) \\
=&(1-\gamma)\sum_{k=0}^{\infty}\gamma^{k}\mathbb{P}_{P^{\text{aug}}}(x_{k}=x, i_{k}=i | x_0=x^0, i_0=i^0,\mu) - (1-\gamma)\mathbf{1}\{x_0=x, i_0=i\} \\
=&d^{\mu}_{P^{\text{aug}},\gamma}(x,i) - (1-\gamma)\mathbf{1}\{x_0=x, i_0=i\}.
\end{aligned} \tag{7}$$

Based on the definition of occupation measures, one can easily write

$$\mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k\sim\mu, x_0=x^0, i_0=i^0\right) = \sum_{(x',i')\in\bar{\mathcal{X}}} d^{\mu}_{P^{\text{aug}},\gamma}(x',i')\sum_{a'\in\bar{\mathcal{A}}}\mu(a'|x',i')r_\lambda(x',i',a').$$

Define the transition probability matrix P^{aug}_{μ} at the $\{(x,i),(x',i')\}$ as

$$P^{\text{aug}}_{\mu}((x,i),(x',i')) := P^{\text{aug}}(x',i'|x,i,\mu(x,i)),$$

and notice that

$$(I-\gamma P^{\text{aug}}_{\mu})\sum_{k=0}^{\infty}\left(\gamma P^{\text{aug}}_{\mu}\right)^k = \sum_{k=0}^{\infty}\left(\gamma P^{\text{aug}}_{\mu}\right)^k - \left(\gamma P^{\text{aug}}_{\mu}\right)^{k+1} = I - \lim_{k\to\infty}\left(\gamma P^{\text{aug}}_{\mu}\right)^k = I.$$

Therefore, the above equation implies $(I-\gamma P^{\text{aug}}_{\mu})^{-1} = \sum_{t=0}^{\infty}\left(\gamma P^{\text{aug}}_{\mu}\right)^k < \infty$. The series is summable because $\sigma(\gamma P^{\text{aug}}_{\mu}) < 1$. Furthermore, one obtains

$$\left(I-\gamma P^{\text{aug}}_{\mu}\right)^{-1}((x,i),(x',i')) = \sum_{k=0}^{\infty}\gamma^k\mathbb{P}(x_k=x', i_k=i' | x_0=x, i_0=i,\mu).$$

Define $d^{\mu}_{P^{\text{aug}},\gamma}$ as the vector $\{d^{\mu}_{P^{\text{aug}},\gamma}(x,i)\}_{(x,i)\in\bar{\mathcal{X}}}$, expression (7) can be rewritten as follows:

$$\left(I-\gamma P^{\text{aug}}_{\mu}\right)^{\top} d^{\mu}_{P^{\text{aug}},\gamma} = \{\mathbf{1}\{x_0=x, i_0=i\}\}_{(x,i)\in\bar{\mathcal{X}}}.$$

On the other hand, by defining $\widehat{P}^{\text{aug}}_{\mu}((x,i),(x',i')) := \widehat{P}^{\text{aug}}(x',i'|x,i,\mu(x,i))$, we can also write

$$\left(I-\gamma \widehat{P}^{\text{aug}}_{\mu}\right)^{\top} d_{\widehat{P}^{\text{aug}},\mu} = \{\mathbf{1}\{x_0=x, i_0=i\}\}_{(x,i)\in\bar{\mathcal{X}}}.$$

Combining the above expressions implies

$$d_{P^{\text{aug}},\mu} - d_{\widehat{P}^{\text{aug}},\mu} - \gamma\left((P^{\text{aug}}_{\mu})^{\top} d_{P^{\text{aug}},\mu} - (\widehat{P}^{\text{aug}}_{\mu})^{\top} d_{\widehat{P}^{\text{aug}},\mu}\right) = 0,$$

which further implies

$$\begin{aligned}
&\left(I-\gamma P^{\text{aug}}_{\mu}\right)^{\top}\left(d_{P^{\text{aug}},\mu} - d_{\widehat{P}^{\text{aug}},\mu}\right) = \gamma\left((P^{\text{aug}}_{\mu})^{\top} - (\widehat{P}^{\text{aug}}_{\mu})^{\top}\right) d_{\widehat{P}^{\text{aug}},\mu} \\
\implies &\left(d_{P^{\text{aug}},\mu} - d_{\widehat{P}^{\text{aug}},\mu}\right) = \left(I-\gamma P^{\text{aug}}_{\mu}\right)^{-\top}\gamma\left((P^{\text{aug}}_{\mu})^{\top} - (\widehat{P}^{\text{aug}}_{\mu})^{\top}\right) d_{\widehat{P}^{\text{aug}},\mu}.
\end{aligned}$$

For any policy μ, one notices that

$$\begin{aligned}
&(1-\gamma)\left(\mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\lambda}(x_k,i_k,a_k) \mid a_k\sim\mu, P^{\text{aug}}_0\right) - \mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\lambda}(x_k,i_k,a_k) \mid a_k\sim\mu, P^{\text{aug}}_0\right)\right) \\
=&\sum_{x^0,i^0} P^{\text{aug}}_0(x^0,i^0) \sum_{(x',i')\in\bar{\mathcal{X}}}\left(d_{P^{\text{aug}},\mu}(x',i') - d_{\widehat{P}^{\text{aug}},\mu}(x',i')\right)\sum_{a'\in\bar{\mathcal{A}}}\mu(a'|x',i')r_{\lambda}(x',i',a') \\
=&\sum_{x^0,i^0} P^{\text{aug}}_0(x^0,i^0) r^{T}_{\lambda,\mu}\left(I-\gamma P^{\text{aug}}_{\mu}\right)^{-\top}\gamma\left((P^{\text{aug}}_{\mu})^{\top} - (\widehat{P}^{\text{aug}}_{\mu})^{\top}\right) d_{\widehat{P}^{\text{aug}},\mu}
\end{aligned}$$

where $r_{\lambda,\mu}(x,i) = r_{\lambda}(x,i,\mu(x,i))$ is the reward functions induced by policy μ. Furthermore, the following expressions can be easily obtained by using the properties of dual MDPs:

$$\begin{aligned}
\left\{\left(I-\gamma P^{\text{aug}}_{\mu}\right)^{-1} r_{\lambda,\mu}\right\}(x,i) =& \sum_{(x',i')\in\bar{\mathcal{X}}}\sum_{k=0}^{\infty}\gamma^k \mathbb{P}_{P^{\text{aug}}}(x_k=x', i_k=i'|x_0=x,i_0=i,\mu) r_{\lambda,\mu}(x',i'), \\
=&\mathbb{E}_{P^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k r_{\lambda,\mu}(x_k,i_k)|x_0=x,i_0=i,\mu\right],
\end{aligned}$$

and

$$\begin{aligned}
&\left\{\left((P^{\text{aug}}_{\mu})^{\top} - (\widehat{P}^{\text{aug}}_{\mu})^{\top}\right) d_{\widehat{P}^{\text{aug}},\mu}\right\}(x',i') \\
=&\sum_{(x,i)\in\bar{\mathcal{X}}}\sum_{k=0}^{\infty}\gamma^k(1-\gamma)\mathbb{P}_{\widehat{P}^{\text{aug}}}(x_k=x,i_k=i|x_0=x^0,i_0=i^0,\mu)\sum_{a\in\bar{\mathcal{A}}}\left(P^{\text{aug}}(x',i'|x,i,a) - \widehat{P}^{\text{aug}}(x',i'|x,i,a)\right)\mu(a|x,i) \\
=&\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k(1-\gamma)\left(P^{\text{aug}}_{\mu}(x',i'|x_k,i_k) - \widehat{P}^{\text{aug}}_{\mu}(x',i'|x_k,i_k)\right)|x_0=x^0,i_0=i^0,\mu\right].
\end{aligned}$$

By using the above results, the following expression holds:

$$\begin{aligned}
&\mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\lambda}(x_k,i_k,a_k) \mid a_k\sim\mu, P^{\text{aug}}_0\right) - \mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\lambda}(x_k,i_k,a_k) \mid a_k\sim\mu, P^{\text{aug}}_0\right) \\
=&\sum_{x^0,i^0} P^{\text{aug}}_0(x^0,i^0)\Bigg\{\gamma\sum_{(x,i)\in\bar{\mathcal{X}}}\mathbb{E}_{P^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k r_{\lambda,\mu}(x_k,i_k)|x_0=x,i_0=i,\mu\right]\cdot \\
&\qquad\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k\left(P^{\text{aug}}_{\mu}(x,i|x_k,i_k) - \widehat{P}^{\text{aug}}_{\mu}(x,i|x_k,i_k)\right)|x_0=x^0,i_0=i^0,\mu\right] \\
\geq&-\frac{2\gamma\lambda R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)\mathbf{1}\{i_k=1\}|a_k\sim\mu, P^{\text{aug}}_0\right]
\end{aligned} \tag{8}$$

where $2\lambda R_{\max}$ is the upper bound for reward function $r_\lambda(x,i,a)$ at $i=1$ with $\lambda \geq 0$.

Thus, by noticing that

$$\begin{aligned}&\mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_\lambda(x_k,i_k,a_k) \mid a_k\sim\mu, P_0^{\text{aug}}\right) - \frac{2\gamma\lambda R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)\mathbf{1}\{i_k=1\}|a_k\sim\mu, P_0^{\text{aug}}\right]\\ =&\mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k \widehat{r}_\lambda(x_k,i_k,a_k) \mid a_k\sim\mu, P_0^{\text{aug}}\right),\end{aligned} \tag{9}$$

and recalling the definition of $\widehat{L}(\mu,\lambda)$, one concludes that $L(\mu,\lambda)-\widehat{L}(\mu,\lambda)\geq 0$. □

By solving for $\max_\mu \min_{\lambda\geq 0}\widehat{L}(\mu,\lambda)$ and guaranteeing that it is lower bounded, one can ensure feasibility of the risk neutral maximization problem.

7 Algorithm

We first state the following algorithm for finding the saddle point $(\widehat{\mu},\widehat{\lambda})$ of the max-min optimization problem $\max_\mu \min_{\lambda\geq 0}\widehat{L}(\mu,\lambda)$. The convergence analysis of this algorithm is given in the following theorem.

Theorem 7 (Convergence)**.** *The policy iteration algorithm terminates in a finite number of steps. Furthermore, the Markov stationary control policy $\widehat{\mu}:=\widehat{\mu}_{\widehat{\lambda}}$ is optimal to the risk neutral optimization problem*

$$\begin{aligned}\max_{\mu\in\Pi_S}\quad & \mathbb{E}_{\widehat{P}}\left(\sum_{k=0}^{\infty}\gamma^k r(x_k,a_k) \mid a_k\sim\mu, P_0\right)\\ \textit{subject to}\quad & \mathbb{E}_{\widehat{P}}\left(\sum_{k=0}^{\infty}\gamma^k r(x_k,a_k) \mid a_k\sim\mu, P_0\right)\geq M_B,\end{aligned}$$

when this algorithm terminates and

$$\sum_{x,i}P_0^{aug}(x,i)\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)-\widehat{\lambda}M_B=\widehat{\mathcal{P}}(P_0^{aug},M_B),$$

where

$$\widehat{\mathcal{P}}(P_0^{aug},M_B)=\max_{\mu}\min_{\lambda\geq 0}\mathbb{E}_{\widehat{P}^{aug}}\left(\sum_{k=0}^{\infty}\gamma^k\widehat{r}_\lambda(x_k,i_k,a_k)\mid a_k\sim\mu,P_0^{aug}\right)-\lambda M_B=\max_{\mu}\min_{\lambda\geq 0}\widehat{L}(\mu,\lambda).$$

Proof. Define the Bellman operators with respect to $\widehat{P}^{\text{aug}}$ and $\widehat{r}_\lambda$ as follows:

$$\widehat{T}_\lambda[V](x,i)=\max_{a\in\bar{\mathcal{A}}}\left\{\widehat{r}_\lambda(x,i,a)+\gamma\sum_{(x',i')\in\bar{\mathcal{X}}}\widehat{P}^{\text{aug}}(x',i'|x,i,a)V(x',i')\right\},$$

and

$$\widehat{T}_\lambda^\mu[V](x,i)=\widehat{r}_\lambda(x,i,\mu(x,i))+\gamma\sum_{(x',i')\in\bar{\mathcal{X}}}\widehat{P}^{\text{aug}}(x',i'|x,i,\mu(x,i))V(x',i').$$

From the policy update rule in the above algorithm, one obtains

$$\widehat{T}_{\lambda^{(j)}}^{\mu^{(q+1)}}\left[\widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}\right](x,i)=\widehat{T}_{\lambda^{(j)}}\left[\widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}\right](x,i),\ \forall(x,i)\in\bar{\mathcal{X}}.$$

On the other hand, by the fixed point equation,

$$\widehat{T}_{\lambda^{(j)}}^{\mu^{(q)}}\left[\widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}\right](x,i)=\widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}(x,i)$$

Input: threshold M_B
Initialization: policy parameter $\mu = \mu^{(0)}$, Lagrangian parameter $\lambda = \lambda^{(0)}$ and initial dual function estimate $f_{\min}^{(0)}$
for $j = 0, 1, 2, \ldots$ **do**
$q \leftarrow 0$.
while TRUE **do**
Policy evaluation: For given policy $\mu^{(q)}$, compute the value function $\widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}$ by solving the following optimization problem,

$$\min_{V} \quad \sum_{x,i} P_0^{\text{aug}}(x,i)V(x,i)$$
$$\text{s.t.} \quad \widehat{r}_{\lambda^{(j)}}(x,i,\mu^{(q)}(x,i)) + \gamma \sum_{(x',i')\in\bar{\mathcal{X}}} \widehat{P}(x'|x,\mu^{(q)}(x,i))V(x',i')\mathbf{1}\{i=i'\} \le V(x,i),\ \forall (x,i)\in\bar{\mathcal{X}}.$$

Policy improvement:

$$\mu^{(q+1)}(x,i) \in \operatorname*{argmax}_{a\in\mathcal{A}} \left\{ \widehat{r}_{\lambda^{(j)}}(x,i,a) + \sum_{(x',i')\in\bar{\mathcal{X}}} \widehat{P}(x'|x,a)\widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}(x',i')\mathbf{1}\{i=i'\} \le \widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}(x,i) \right\}.$$

$q \leftarrow q+1$.
if $\widehat{V}_{\lambda^{(j)}}^{\mu^{(q)}}(x,i) = \widehat{V}_{\lambda^{(j)}}^{\mu^{(q+1)}}(x,i),\ \forall (x,i)\in\bar{\mathcal{X}}$ **then**
return $\widehat{\mu}_{\lambda^{(j)}} = \mu^{(q)}$ and **break**.
Lagrange multiplier update:

$$\lambda^{(j+1)} = \left(\lambda^{(j)} - \alpha^{(j)} \left(\mathbb{E}_{\widehat{P}} \left(\sum_{k=0}^{\infty} \gamma^k \left(r(x_k,a_k) - \frac{\gamma R_{\max}}{1-\gamma} e(x_k,a_k) \right) \mid a_k \sim \widehat{\mu}_{\lambda^{(j)}}, P_0 \right) - M_B \right) \right)^+$$

where the step length $\alpha^{(j)}$ satisfies the following conditions:

$$\alpha^{(j)} \ge 0,\ \sum_{j=0}^{\infty} \alpha^{(j)} = \infty,\ \sum_{j=0}^{\infty} \left(\alpha^{(j)}\right)^2 < \infty. \tag{10}$$

Best dual function estimate update: $f_{\min}^{(j+1)} = \min\left(f_{\min}^{(j)}, \widehat{L}(\widehat{\mu}_{\lambda^{(j+1)}}, \lambda^{(j+1)}) \right)$.
Best Lagrange multiplier update:

$$\lambda^{(j+1)} \leftarrow \begin{cases} \lambda^{(j+1)} & \text{if } f_{\min}^{(j+1)} = \widehat{L}(\widehat{\mu}_{\lambda^{(j+1)}}, \lambda^{(j+1)}) \\ \lambda^{(j)} & \text{otherwise} \end{cases}.$$

Terminate algorithm when $\lambda^{(j)}$ converges to $\widehat{\lambda}$.

and notice that $\widehat{T}^{\mu^{(q)}}_{\lambda^{(j)}}[V](x,i) \leq \widehat{T}_{\lambda^{(j)}}[V](x,i)$, one obtains

$$\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}(x,i) \leq \widehat{T}^{\mu^{(q+1)}}_{\lambda^{(j)}}\left[\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}\right](x,i).$$

By repeatedly applying $\widehat{T}^{\mu^{(q+1)}}_{\lambda^{(j)}}$ on both sides, this inequality also implies

$$\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}(x,i) \leq \widehat{T}^{\mu^{(q+1)}}_{\lambda^{(j)}}\left[\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}\right](x,i) \leq \cdots \leq \lim_{n\to\infty}\left(\widehat{T}^{\mu^{(q)}}_{\lambda^{(j)}}\right)^n\left[\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}\right](x,i) \leq \widehat{V}^{\mu^{(q+1)}}_{\lambda^{(j)}}(x,i) \quad (11)$$

for any $(x,i) \in \bar{\mathcal{X}}$. When the inner-loop stopping condition is satisfied: $\widehat{V}^{\mu^{(q+1)}}_{\lambda^{(j)}}(x,i) = \widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}(x,i)$, then

$$\widehat{T}_{\lambda^{(j)}}\left[\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}\right](x,i) = \widehat{T}^{\mu^{(q+1)}}_{\lambda^{(j)}}\left[\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}\right](x,i) = \widehat{T}^{\mu^{(q+1)}}_{\lambda^{(j)}}\left[\widehat{V}^{\mu^{(q+1)}}_{\lambda^{(j)}}\right](x) = \widehat{V}^{\mu^{(q+1)}}_{\lambda^{(j)}}(x,i) = \widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}(x,i)$$

The first equality is based on the definition of Bellman operator $\widehat{T}_\lambda$ and the policy iteration algorithm. The third equality is due to the fact that $\widehat{V}^{\mu^{(q+1)}}_{\lambda^{(j)}}$ is a solution to fixed point equation: $\widehat{T}^{\mu^{(q+1)}}_{\lambda^{(j)}}[V](x,i) = V(x,i)$. The second and the fourth equality are due to the inner-loop stopping condition. Notice that the fixed point of $\widehat{T}_{\lambda^{(j)}}[V](x,i) = V(x,i)$ is unique. Combining these arguments, one obtains $\widehat{V}^{\mu^{(q)}}_{\lambda^{(j)}}(x,i) = \widehat{V}_{\lambda^{(j)}}(x,i)$ where $\widehat{V}_{\lambda^{(j)}}(x,i) = \max_\mu \mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^\infty \gamma^k \widehat{r}_{\lambda^{(j)}}(x_k,i_k,a_k) \mid a_k \sim \mu, x_0 = x, i_0 = i\right)$. This implies $\widehat{\mu}_{\lambda^{(j)}} = \mu^{(q)} \in \arg\max_\mu \widehat{L}(\mu,\lambda^{(j)})$.

Second we want to show that the sequence of $\lambda^{(j)}$ converges to the global minimum $\widehat{\lambda} \in \operatorname{argmin}_{\lambda\geq 0} f(\lambda)$ where $f(\lambda) = \max_\mu \widehat{L}(\mu,\lambda)$. Notice by definition that $\widehat{L}(\mu,\lambda)$ is a linear function of λ. Furthermore, since the policy μ belongs to the closed subset of the compact set of simplexes, then $\max_\mu \widehat{L}(\mu,\lambda)$ is the point-wise supremum over an infinite set of convex functions. This implies that $f(\lambda) = \max_\mu \widehat{L}(\mu,\lambda)$ is convex in λ. It can be easily shown by sub-gradient calculus that

$$\left.\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\right|_{\mu=\mu^*_\lambda} = \mathbb{E}_{\widehat{P}}\left(\sum_{k=0}^\infty \gamma^k\left(r(x_k,a_k) - \frac{\gamma R_{\max}}{1-\gamma}e(x_k,a_k)\right) \mid a_k \sim \widehat{\mu}_\lambda, P_0\right) - M_B \in \partial f(\lambda).$$

This implies that the Lagrange multiplier update can be viewed as projected sub-gradient descent. Now, for

$$\bar{\lambda}^{(j+1)} := \lambda^{(j)} - \alpha^{(j)}\left.\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\right|_{\mu=\widehat{\mu}_\lambda,\lambda=\lambda^{(j)}}$$

and $\lambda^{(j+1)} = \left(\bar{\lambda}^{(j+1)}\right)^+$, since the projection operator is non-expansive, one obtains $(\lambda^{(j+1)} - \widehat{\lambda})^2 \leq (\bar{\lambda}^{(j+1)} - \widehat{\lambda})^2$.

Therefore, the following expression holds:

$$\begin{aligned}
(\lambda^{(j+1)} - \widehat{\lambda})^2 \leq & (\bar{\lambda}^{(j+1)} - \widehat{\lambda})^2 \\
= & \left(\lambda^{(j)} - \alpha^{(j)}\left.\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\right|_{\mu=\widehat{\mu}_\lambda,\lambda=\lambda^{(j)}} - \widehat{\lambda}\right)^2 \\
= & (\lambda^{(j)} - \widehat{\lambda})^2 - 2\alpha^{(j)}(\lambda^{(j)} - \widehat{\lambda})\left.\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\right|_{\mu=\widehat{\mu}_\lambda,\lambda=\lambda^{(j)}} + \left(\alpha^{(j)}\right)^2\left(\left.\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\right|_{\mu=\widehat{\mu}_\lambda,\lambda=\lambda^{(j)}}\right)^2 \\
\leq & (\lambda^{(j)} - \widehat{\lambda})^2 - 2\alpha^{(j)}(f(\lambda^{(j)}) - f(\widehat{\lambda})) + \left(\alpha^{(j)}\right)^2\left(\left.\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\right|_{\mu=\widehat{\mu}_\lambda,\lambda=\lambda^{(j)}}\right)^2,
\end{aligned}$$

The inequality is due to the fact that for any $\lambda \geq 0$ and $d\widehat{L}(\mu,\lambda)/d\lambda|_{\mu=\widehat{\mu}_\lambda} \in \partial f(\lambda)$,

$$f(\lambda) - f(\widehat{\lambda}) \geq (\lambda - \widehat{\lambda})\left.\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\right|_{\mu=\widehat{\mu}_\lambda}.$$

This further implies

$$(\lambda^{(j+1)} - \widehat{\lambda})^2 \leq (\lambda^{(0)} - \widehat{\lambda})^2 - \sum_{q=0}^{j} 2\alpha^{(q)}(f(\lambda^{(q)}) - f(\widehat{\lambda})) + \left(\alpha^{(q)}\right)^2 \left(\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\Big|_{\mu=\widehat{\mu}_\lambda, \lambda=\lambda^{(q)}} \right)^2.$$

Since $(\lambda^{(j+1)} - \widehat{\lambda})^2$ is a positive quantity and $(\lambda^{(0)} - \widehat{\lambda})^2$ is bounded, this further implies

$$2\sum_{q=0}^{j} \alpha^{(q)}(f(\lambda^{(q)}) - f(\widehat{\lambda})) \leq (\bar{\lambda}^{(0)} - \widehat{\lambda})^2 + \sum_{q=0}^{j} \left(\alpha^{(q)}\right)^2 \left(\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\Big|_{\mu=\widehat{\mu}_\lambda, \lambda=\lambda^{(q)}} \right)^2.$$

By defining $f_{\text{min}}^{(j)} = \min_{q\in\{0,\ldots,j\}} f(\lambda^{(q)})$, the above expression implies

$$f_{\text{min}}^{(j)} - f(\widehat{\lambda}) \leq \frac{1}{\sum_{q=0}^{j} \alpha^{(q)}} \left((\lambda^{(0)} - \lambda^*)^2 + \sum_{q=0}^{j} \left(\alpha^{(q)}\right)^2 \left(\frac{d\widehat{L}(\mu,\lambda)}{d\lambda}\Big|_{\mu=\widehat{\mu}_\lambda, \lambda=\lambda^{(q)}} \right)^2 \right).$$

The step-size rule in (10) for $\alpha^{(q)}$ ensures that the numerator is bounded and the denominator goes to infinity as $j \to \infty$. This implies that for any $\epsilon > 0$, there exists a constant $N(\epsilon)$ such that for any $j > N(\epsilon)$, $f(\widehat{\lambda}) \leq f_{\text{min}}^{(j)} \leq f(\widehat{\lambda}) + \epsilon$. In other words, the sequence $\lambda^{(j)}$ converges to the global minimum $\widehat{\lambda}$ of $f(\lambda)$ where $f(\lambda) = \max_\mu \widehat{L}(\mu, \lambda)$.

By combining all previous arguments, one shows that $\widehat{L}(\widehat{\mu}, \widehat{\lambda}) = \widehat{\mathcal{D}}(P_0^{\text{aug}}, M_B)$, where

$$\widehat{\mathcal{D}}(P_0^{\text{aug}}, M_B) = \min_{\lambda\geq 0} \max_{\mu} \mathbb{E}_{\widehat{P}^{\text{aug}}} \left(\sum_{k=0}^{\infty} \gamma^k \widehat{r}_\lambda(x_k, i_k, a_k) \mid a_k \sim \mu, P_0^{\text{aug}} \right) - \lambda M_B = \min_{\lambda\geq 0} \max_{\mu} \widehat{L}(\mu, \lambda).$$

Then the convergence proof is completed by using strong duality, i.e., $\widehat{\mathcal{D}}(P_0^{\text{aug}}, M_B) = \widehat{\mathcal{P}}(P_0^{\text{aug}}, M_B)$. □

In the next section, we will investigate the performance in terms of sub-optimality for the solution of this algorithm.

8 Performance

We are now in position to derive a performance bound on the control policies found by the above algorithm. By the primal formulation of Lagrangian function, one obtains the following expressions:

$$\begin{aligned} \max_{\mu} \min_{\lambda\geq 0} L(\mu, \lambda) = \max_{\mu\in\Pi_S} \quad & \mathbb{E}_{\widehat{P}} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu, P_0 \right) \\ \text{subject to} \quad & \mathbb{E}_{P} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu, P_0 \right) \geq M_B, \end{aligned}$$

and $\min_{\lambda\geq 0} L(\widehat{\mu}, \lambda) = \mathbb{E}_{\widehat{P}} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \widehat{\mu}, P_0 \right)$ such that $\mathbb{E}_{P} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \widehat{\mu}, P_0 \right) \geq M_B$. Let μ^* be the optimal policy of the above constrained problem. In order to compare the performance of the policy from approximation $\widehat{\mu}$, one needs to calculate the upper/lower bound for the following expression:

$$\mathbb{E}_{P} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu^*, P_0 \right) - \mathbb{E}_{P} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \widehat{\mu}, P_0 \right).$$

By the primal formulation of Lagrangian function, since both $\widehat{\mu}$ and μ^* are feasible control policies and the stage-wise reward functions are bounded, one can easily check that the following quantities are bounded: $\min_{\lambda\geq 0} L(\widehat{\mu}, \lambda)$ and $\max_\mu \min_{\lambda\geq 0} L(\mu, \lambda)$. Furthermore, this also implies that $\mathbb{E}_{P} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu^*, P_0 \right) - \mathbb{E}_{P} \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \widehat{\mu}, P_0 \right) = \max_\mu \min_{\lambda\geq 0} L(\mu, \lambda) - \min_{\lambda\geq 0} L(\widehat{\mu}, \lambda)$.

Define the Bellman residual as follows:

$$\text{BR}(\widehat{\mu},\widehat{\lambda}) = \left\| T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\widehat{\mu}}] - V_{\widehat{\lambda}}^{\widehat{\mu}} \right\|_{\infty}.$$

Based on the definitions of stage-wise reward functions $r_{\widehat{\lambda}}$ and $\widehat{r}_{\widehat{\lambda}}$, when the probability mis-match error function $e(x,i,a)$ tends to zero for any (x,i,a), the fixed point solution calculated by the approximated policy $V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)$ converges to the fixed point solution induced by the optimal policy $V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)$ for any (x,i). Notice that $T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\tilde{\mu}}](x,i) = V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)$ for any (x,i). Therefore, when $\max_{x,i,a} e(x,i,a) \rightarrow 0$, one obtains $V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) \rightarrow V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)$, $T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i) \rightarrow T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\tilde{\mu}}](x,i) = V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)$ for any (x,i) and the Bellman residual $\text{BR}(\widehat{\mu},\widehat{\lambda})$ tends to zero.

First, we introduce the following Lemma on a sub-optimality performance guarantee.

Lemma 8 (Performance)**.** *Let $(\widehat{\mu},\widehat{\lambda})$ be the saddle point solution of $\min_{\lambda\geq 0}\max_{\mu}\widehat{L}(\mu,\lambda)$. Then the following performance bound holds:*

$$\begin{aligned}
0 \leq & \mathbb{E}_P\left(\sum_{k=0}^{\infty}\gamma^k r(x_k,a_k) \mid a_k \sim \mu^*, P_0\right) - \mathbb{E}_P\left(\sum_{k=0}^{\infty}\gamma^k r(x_k,a_k) \mid a_k \sim \widehat{\mu}, P_0\right) \\
\leq & \frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)|a_k\sim\widehat{\mu},P_0\right] + \frac{\textit{BR}(\widehat{\mu},\widehat{\lambda})}{1-\gamma}.
\end{aligned}$$

Proof. Recall that $\widehat{L}(\mu,\lambda) \leq L(\mu,\lambda)$ for any μ and $\lambda \geq 0$. Therefore, by defining $\tilde{\mu} \in \arg\max_{\mu} L(\mu,\widehat{\lambda})$, the following expression holds:

$$\begin{aligned}
0 \leq & \min_{\lambda\geq 0}\max_{\mu} L(\mu,\lambda) - \min_{\lambda\geq 0}\max_{\mu}\widehat{L}(\mu,\lambda) \leq \max_{\mu} L(\mu,\widehat{\lambda}) - \widehat{L}(\widehat{\mu},\widehat{\lambda}) = L(\tilde{\mu},\widehat{\lambda}) - \widehat{L}(\widehat{\mu},\widehat{\lambda}) \\
= & \sum_{x,i} P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i) - \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right) \\
= & \sum_{x,i} P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i) - V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) + V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) - \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right)
\end{aligned} \tag{12}$$

The first inequality is due to the fact that $\widehat{L}(\mu,\lambda) \leq L(\mu,\lambda)$. The second inequality is based on the definition of the saddle point solution of $\min_{\lambda\geq 0}\max_{\mu}\widehat{L}(\mu,\lambda)$, and the fact that $\widehat{\lambda}$ is a feasible solution to $\min_{\lambda\geq 0}\max_{\mu} L(\mu,\lambda)$. The first equality follows from definition and the second equality follows from the result of Bellman's equation. The last equality is a result of basic algebraic manipulations.

In order to derive a performance bound between the approximated solution and the optimal solution. We study the rightmost term $\sum_{x,i} P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i) - V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) + V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) - \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right)$ of expression (12). Consider the first part of this expression:

$$\sum_{x,i} P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) - \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right).$$

Based on previous analysis in Bellman's recursion, one can easily see that

$$\begin{aligned}
\sum_{x,i} P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) - \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right) = & \mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k\left(r_{\widehat{\lambda}}(x_k,i_k,a_k) - \widehat{r}_{\widehat{\lambda}}(x_k,i_k,a_k)\right) \mid a_k\sim\widehat{\mu}, P_0^{\text{aug}}\right) \\
& + \mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\widehat{\lambda}}(x_k,i_k,a_k) \mid a_k\sim\widehat{\mu}, P_0^{\text{aug}}\right) - \mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\widehat{\lambda}}(x_k,i_k,a_k) \mid a_k\sim\widehat{\mu}, P_0^{\text{aug}}\right).
\end{aligned}$$

By expression (9), one notices that

$$\begin{aligned}
& \mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k(r_{\widehat{\lambda}}(x_k,i_k,a_k) - \widehat{r}_{\widehat{\lambda}}(x_k,i_k,a_k)) \mid a_k\sim\widehat{\mu}, P_0^{\text{aug}}\right) \\
= & \frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)\mathbf{1}\{i_k=1\}|a_k\sim\widehat{\mu},P_0^{\text{aug}}\right] = \frac{\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)|a_k\sim\widehat{\mu},P_0\right].
\end{aligned}$$

On the other hand, similar to the analysis in expression (8), one obtains

$$
\begin{aligned}
&\mathbb{E}_{P^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\widehat{\lambda}}(x_k,i_k,a_k)\mid a_k\sim\widehat{\mu},P_0^{\text{aug}}\right)-\mathbb{E}_{\widehat{P}^{\text{aug}}}\left(\sum_{k=0}^{\infty}\gamma^k r_{\widehat{\lambda}}(x_k,i_k,a_k)\mid a_k\sim\widehat{\mu},P_0^{\text{aug}}\right)\\
=&\sum_{x^0,i^0}P_0^{\text{aug}}(x^0,i^0)\Bigg\{\gamma\sum_{(x,i)\in\bar{\mathcal{X}}}\mathbb{E}_{P^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k r_{\widehat{\lambda},\mu}(x_k,i_k)|x_0=x,i_0=i,a_k\sim\widehat{\mu}\right]\cdot\\
&\qquad\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k\left(P_\mu^{\text{aug}}(x,i|x_k,i_k)-\widehat{P}_\mu^{\text{aug}}(x,i|x_k,i_k)\right)|x_0=x^0,i_0=i^0,a_k\sim\widehat{\mu}\right]\\
\leq&\frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)\mathbf{1}\{i_k=1\}|a_k\sim\widehat{\mu},P_0^{\text{aug}}\right]\\
=&\frac{\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)|a_k\sim\widehat{\mu},P_0\right].
\end{aligned}
\tag{13}
$$

Therefore by combining these analysis,

$$
\sum_{x,i}P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)-\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right)\leq\frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)|a_k\sim\widehat{\mu},P_0\right]. \tag{14}
$$

Next consider the second part in expression (12):

$$
\sum_{x,i}P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)-V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right).
$$

Recall the Bellman residual: $\text{BR}(\widehat{\mu},\widehat{\lambda})=\left\|T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\widehat{\mu}}]-V_{\widehat{\lambda}}^{\widehat{\mu}}\right\|_\infty$. By applying the contraction mapping property on $T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)-V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)$ for any (x,i), one obtains

$$
T_{\widehat{\lambda}}^2[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)-T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)\leq\gamma\text{BR}(\widehat{\mu},\widehat{\lambda}).
$$

By an induction argument, the above expression becomes $T_{\widehat{\lambda}}^N[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)-T_{\widehat{\lambda}}^{N-1}[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)\leq\gamma^{N-1}\text{BR}(\widehat{\mu},\widehat{\lambda})$, for which by a telescoping sum, it further implies

$$
T_{\widehat{\lambda}}^N[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)-V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)=\sum_{k=1}^{N}T_{\widehat{\lambda}}^k[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)-T_{\widehat{\lambda}}^{k-1}[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)\leq\sum_{k=1}^{N}\gamma^{k-1}\text{BR}(\widehat{\mu},\widehat{\lambda}).
$$

By letting $N\to\infty$ and noticing that $\lim_{N\to\infty}T_{\widehat{\lambda}}^N[V_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)=V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)$, one finally obtains

$$
\sum_{x,i}P_0^{\text{aug}}(x,i)\left(V_{\widehat{\lambda}}^{\tilde{\mu}}(x,i)-V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right)\leq\lim_{N\to\infty}\sum_{k=1}^{N}\gamma^{k-1}\text{BR}(\widehat{\mu},\widehat{\lambda})=\frac{\text{BR}(\widehat{\mu},\widehat{\lambda})}{1-\gamma}.
$$

By combining both parts in expression (12), the performance bound can be re-written as follows:

$$
\begin{aligned}
0\leq\min_{\lambda\geq0}\max_{\mu}L(\mu,\lambda)-\min_{\lambda\geq0}\max_{\mu}\widehat{L}(\mu,\lambda)\leq&\frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)|a_k\sim\widehat{\mu},P_0\right]\\
&+\frac{\text{BR}(\widehat{\mu},\widehat{\lambda})}{1-\gamma}.
\end{aligned}
\tag{15}
$$

Furthermore, by defining $\tilde{\lambda} \in \arg\min_{\lambda\geq 0} L(\widehat{\mu},\lambda)$, one obtains

$$\begin{aligned}
0 \leq & \max_{\mu}\min_{\lambda\geq 0} L(\mu,\lambda) - L(\widehat{\mu},\tilde{\lambda}) \\
\leq & \min_{\lambda\geq 0}\max_{\mu} L(\mu,\lambda) - \widehat{L}(\widehat{\mu},\tilde{\lambda}) \\
\leq & \min_{\lambda\geq 0}\max_{\mu} L(\mu,\lambda) - \min_{\lambda\geq 0}\widehat{L}(\widehat{\mu},\lambda) \\
= & \min_{\lambda\geq 0}\max_{\mu} L(\mu,\lambda) - \max_{\mu}\min_{\lambda\geq 0}\widehat{L}(\mu,\lambda) \\
= & \min_{\lambda\geq 0}\max_{\mu} L(\mu,\lambda) - \min_{\lambda\geq 0}\max_{\mu}\widehat{L}(\mu,\lambda) \\
\leq & \frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}^{\text{aug}}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)\mathbf{1}\{i_k=1\}|a_k\sim\widehat{\mu},P_0\right] + \frac{\text{BR}(\widehat{\mu},\widehat{\lambda})}{1-\gamma}.
\end{aligned}$$

The first inequality follows from the fact that $L(\widehat{\mu},\tilde{\lambda}) = \min_{\lambda\geq 0} L(\widehat{\mu},\lambda) \leq \max_{\mu}\min_{\lambda\geq 0} L(\mu,\lambda)$, where the first and second inequality in this expression follow from definition and the last equality follows from strong duality. The second inequality is based on the fact that $L(\mu,\lambda) \geq \widehat{L}(\mu,\lambda)$ and strong duality $\max_{\mu}\min_{\lambda\geq 0} L(\mu,\lambda) = \min_{\lambda\geq 0}\max_{\mu} L(\mu,\lambda)$. The third inequality and the first equality follows from definition. The second equality follows from strong duality: $\min_{\lambda\geq 0}\max_{\mu}\widehat{L}(\mu,\lambda) = \max_{\mu}\min_{\lambda\geq 0}\widehat{L}(\mu,\lambda)$ and the last inequality follows from the performance bound in (15). □

Recall the Bellman residual: $\text{BR}(\widehat{\mu},\widehat{\lambda}) = \left\|T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\widehat{\mu}}] - V_{\widehat{\lambda}}^{\widehat{\mu}}\right\|_{\infty}$. In general, one cannot calculate the Bellman residual $\text{BR}(\widehat{\mu},\widehat{\lambda})$ because the transition probability P^{aug} is not known. We therefore provide an upper bound for $\text{BR}(\widehat{\mu},\widehat{\lambda})$ in the next lemma.

Lemma 9. *Let $(\widehat{\mu},\widehat{\lambda})$ be the saddle point solution of $\min_{\lambda\geq 0}\max_{\mu}\widehat{L}(\mu,\lambda)$. Then the following expression holds:*

$$BR(\widehat{\mu},\widehat{\lambda}) \leq \max_{x\in\mathcal{X},a\in\mathcal{A}}\left\{\frac{2(1+\gamma)\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)|x_0=x,a_k\sim\widehat{\mu}\right] + \frac{4\gamma\widehat{\lambda}R_{\max}}{1-\gamma}e(x,a)\right\}.$$

Proof. By triangular inequality, one can easily see that

$$\begin{aligned}
\text{BR}(\widehat{\mu},\widehat{\lambda}) \leq & \left\|T_{\widehat{\lambda}}[\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}] - \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}\right\|_{\infty} + \left\|\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}} - V_{\widehat{\lambda}}^{\widehat{\mu}}\right\|_{\infty} + \left\|T_{\widehat{\lambda}}[V_{\widehat{\lambda}}^{\widehat{\mu}}] - T_{\widehat{\lambda}}[\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}]\right\|_{\infty} \\
\leq & (1+\gamma)\left\|\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}} - V_{\widehat{\lambda}}^{\widehat{\mu}}\right\|_{\infty} + \left\|T_{\widehat{\lambda}}[\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}] - \widehat{T}_{\widehat{\lambda}}[\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}]\right\|_{\infty}
\end{aligned}$$

The second inequality is due to contraction property of Bellman operator and the fixed point theorem: $\widehat{T}_{\widehat{\lambda}}[\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}](x,i) = \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)$, for any $(x,i)\in\bar{\mathcal{X}}$. Similar to the derivation in (14), one obtains

$$\left|V_{\widehat{\lambda}}^{\widehat{\mu}}(x,i) - \widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x,i)\right| \leq \frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\mathbb{E}_{\widehat{P}}\left[\sum_{k=0}^{\infty}\gamma^k e(x_k,a_k)|x_0=x,a_k\sim\widehat{\mu}\right], \forall(x,i)\in\bar{\mathcal{X}}.$$

Furthermore for any $(x,i)\in\bar{\mathcal{X}}$,

$$\begin{aligned}
\left|T_{\widehat{\lambda}}[\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}](x,i) - \widehat{T}_{\widehat{\lambda}}[\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}](x,i)\right| = & \left|\max_{a\in\mathcal{A}}\left\{r_{\widehat{\lambda}}(x,i,a) + \gamma\sum_{(x',i')\in\bar{\mathcal{X}}}P^{\text{aug}}(x',i'|x,i,a)\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x',i')\right\}\right. \\
& \left. - \max_{a\in\mathcal{A}}\left\{\widehat{r}_{\widehat{\lambda}}(x,i,a) + \gamma\sum_{(x',i')\in\bar{\mathcal{X}}}\widehat{P}^{\text{aug}}(x',i'|x,i,a)\widehat{V}_{\widehat{\lambda}}^{\widehat{\mu}}(x',i')\right\}\right| \\
\leq & \max_{a\in\mathcal{A}}\left\{\sum_{x'\in\mathcal{X}}\left|P(x'|x,a) - \widehat{P}(x'|x,a)\right|\frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma} + \frac{2\gamma\widehat{\lambda}R_{\max}}{1-\gamma}e(x,a)\right\} \leq \frac{4\gamma\widehat{\lambda}R_{\max}}{1-\gamma}\max_{a\in\mathcal{A}}e(x,a).
\end{aligned}$$

Thus by combining the above results, the proof is completed. □

Finally, we provide the sub-optimality performance bound in the following theorem. The proof of this theorem is completed by combining the results from Lemma 8 and 9.

Theorem 10 (Performance)**.** *Let* $(\widehat{\mu}, \widehat{\lambda})$ *be the saddle point solution of* $\min_{\lambda \geq 0} \max_{\mu} \widehat{L}(\mu, \lambda)$*. Then the following performance bound holds:*

$$
\begin{aligned}
0 \leq & \mathbb{E}_P \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \mu^*, P_0 \right) - \mathbb{E}_P \left(\sum_{k=0}^{\infty} \gamma^k r(x_k, a_k) \mid a_k \sim \widehat{\mu}, P_0 \right) \\
\leq & \frac{2\gamma \widehat{\lambda} R_{\max}}{1-\gamma} \mathbb{E}_{\widehat{P}} \left[\sum_{k=0}^{\infty} \gamma^k e(x_k, a_k) | a_k \sim \widehat{\mu}, P_0 \right] \\
& + \max_{x \in \mathcal{X}, a \in \mathcal{A}} \left\{ \frac{2(1+\gamma)\gamma \widehat{\lambda} R_{\max}}{(1-\gamma)^2} \mathbb{E}_{\widehat{P}} \left[\sum_{k=0}^{\infty} \gamma^k e(x_k, a_k) | x_0 = x, a_k \sim \widehat{\mu} \right] + \frac{4\gamma \widehat{\lambda} R_{\max}}{(1-\gamma)^2} e(x, a) \right\}.
\end{aligned}
$$

9 Conclusions

In this paper, we provided an offline algorithm that solves for a "good" control policy using a simulated Markov decision process (MDP), and guaranteed that this policy performs better than the baseline policy in the real environment. Furthermore, we provided a performance bound on sub-optimality for the control policy generated by this algorithm.

Future work includes 1) an extension of the problem formulation to include risk-sensitive objective functions, 2) an analysis the convergence rate to the current algorithm, 3) numerical experiments of this algorithm in practical domains and 4) a development of approximate algorithms when the state and control spaces are exponentially huge.

References

[1] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.

[2] Jean Pierre Aubin. A Pareto minimum principle. Technical report, DTIC Document, 1972.

[3] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific Belmont, MA, 1995.

[4] John E Goulionis and Athanassios Vozikis. Medical decision making for patients with Parkinson disease under average cost criterion. *Australia and New Zealand health policy*, 6(1):15, 2009.

[5] Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

[6] Laurent Jeanpierre and François Charpillet. Automated medical diagnosis with fuzzy stochastic models: monitoring chronic diseases. *Acta Biotheoretica*, 52(4):291–311, 2004.

[7] Melody Y Kiang, TS Raghu, and Kevin Huei-Min Shang. Marketing on the internet−who can benefit from an online marketing approach? *Decision Support Systems*, 27(4):383–393, 2000.

[8] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670. ACM, 2010.

[9] Babak Mahmoudi, Jose C Principe, and Justin C Sanchez. Extracting an evaluative feedback from the brain for adaptation of motor neuroprosthetic decoders. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, pages 1682–1685. IEEE, 2010.

[10] Arnab Nilim and Laurent El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

[11] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*, volume 414. John Wiley & Sons, 2009.

[12] Kevin Regan and Craig Boutilier. Robust policy computation in reward-uncertain MDPs using nondominated policies. In *AAAI*, 2010.

[13] JA Smith, JS Gaynor, RM Bednarski, and WW Muir. Adverse effects of administration of Propofol with various preanesthetic regimens in dogs. *Journal of the American Veterinary Medical Association*, 202(7): 1111–1115, 1993.

[14] Huan Xu and Shie Mannor. Distributionally robust Markov decision processes. In *Advances in Neural Information Processing Systems*, pages 2505–2513, 2010.